arXiv:1603.08485v1 [cs.CG] 28 Mar 2016

Incremental Voronoi diagrams

Sarah R. Allen* Luis Barba† John Iacono‡ Stefan Langerman§

March 29, 2016

Abstract

We study the amortized number of combinatorial changes (edge insertions and removals) needed to update the graph structure of the Voronoi diagram VD(S) (and several variants thereof) of a set S of n sites in the plane as sites are added to the set. To that effect, we define a general update operation for planar graphs that can be used to model the incremental construction of several variants of Voronoi diagrams as well as the incremental construction of an intersection of halfspaces in $\mathbb{R}^3$. We show that the amortized number of edge insertions and removals needed to add a new site to the Voronoi diagram is $O(\sqrt{n})$. A matching $\Omega(\sqrt{n})$ combinatorial lower bound is shown, even in the case where the graph representing the Voronoi diagram is a tree. This contrasts with the $O(\log n)$ upper bound of Aronov et al. (2006) for farthest-point Voronoi diagrams in the special case where the points are inserted in clockwise order along their convex hull.

We then present a semi-dynamic data structure that maintains the Voronoi diagram of a set S of n sites in convex position. This data structure supports the insertion of a new site p (and hence the addition of its Voronoi cell) and finds the asymptotically minimal number K of edge insertions and removals needed to obtain the diagram of $S \cup \{p\}$ from the diagram of S, in time $O(K \operatorname{polylog} n)$ worst case, which is $O(\sqrt{n} \operatorname{polylog} n)$ amortized by the aforementioned combinatorial result.

The most distinctive feature of this data structure is that the graph of the Voronoi diagram is maintained explicitly at all times and can be retrieved and traversed in the natural way; this contrasts with other known data structures supporting nearest neighbor queries. Our data structure supports general search operations on the current Voronoi diagram, which can, for example, be used to perform point location queries in the cells of the current Voronoi diagram in $O(\log n)$ time, or to determine whether two given sites are neighbors in the Delaunay triangulation.

1 Introduction

Let S be a set of n sites in the plane. The graph structures of the Voronoi diagram VD(S) and its dual the Delaunay triangulation DT(S) capture much of the proximity information of that set. They contain the nearest neighbor graph, the minimum spanning tree, and the Gabriel graph of S, and have countless applications in computational geometry, shape reconstruction, computational biology, and machine learning.

One of the most popular algorithms for constructing a Voronoi diagram inserts sites in random order, incrementally updating the diagram [8]. In that case, backward analysis shows that the expected number of changed edges in VD(S) is constant, offering some hope that an efficient dynamic—or at least semi-dynamic—data structure for maintaining VD(S) could exist. These hopes, however, are rapidly squashed, as it is easy to construct examples where the complexity of each successively added face is $\Omega(n)$, and thus each insertion changes the position of a linear number of vertices and edges of VD(S). The goal of this paper is to show that despite this worst-case behavior, the amortized number of *structural changes* to the graph of

*Carnegie Mellon University. Supported by NSF grants CCF-0747250, CCF-1116594, and the Graduate Research Fellowship Program under Grant No. DGE-1252522. srallen@cs.cmu.edu

†Département d'Informatique, Université Libre de Bruxelles, Carleton University, lbarbafl@ulb.ac.be+

‡Tandon School of Engineering, New York University, iacono@nyu.edu

§Directeur de Recherches du F.R.S.-FNRS, stefan.langerman@ulb.ac.be

the Voronoi diagram of S, that is, the minimum number of edge insertions and deletions needed to update $\mathrm{VD}(S)$ throughout any sequence of site insertions to S, is much smaller.

This might come as a surprise in light of the fact that the number of combinatorial changes (usually modeled as *flips*) to the Delaunay triangulation of S upon the insertion of a point can be $\Omega(n)$ with each insertion, even when the sites are in convex position and are added in clockwise order. (Note that in that case the Voronoi diagram of S is a tree and the standard flip operation is a rotation in the tree.)

To overcome this worst-case behavior, Aronov et al. [2] studied what happens in this specific case (points in convex position added in clockwise order) if the rotation operation is replaced by the more elementary link (add an edge) and cut (delete an edge) operations in the tree. They show that, in that model, it is possible to reconfigure the tree after each site insertion while performing $O(\log n)$ links and cuts, amortized; however their proof is existential and no algorithm is provided to find those links and cuts. Pettie [14] shows both an alternate proof of that fact using forbidden 0-1 matrices and a matching lower bound.

One important application of Voronoi diagrams is to solve *nearest-neighbor* (or *farthest-neighbor*) queries: given a point in the plane, find the site nearest (or farthest) to this point. In the static case, this is done by preprocessing the (nearest or farthest point) Voronoi diagram to answer point-location queries in $O(\log n)$ time. Without the need to maintain $\mathrm{VD}(S)$ explicitly, the problem of nearest neighbor queries is a *decomposable search problem* and can be made semi-dynamic using the standard dynamization techniques of Bentley and Saxe [4]. The best incremental data structure supporting nearest-neighbor queries performs queries and insertions in $O(\log^2 n/\log\log n)$ time [7, 13]. Recently, Chan [6] developed a randomized data structure supporting nearest-neighbor queries in $O(\log^2 n)$ time, insertions in $O(\log^3 n)$ expected amortized time, and deletions in $O(\log^6 n)$ expected amortized time.

Flarbs

In the mid-1980's it was observed that a number of variants of Voronoi diagrams and Delaunay triangulations using different metrics (Euclidean distance, L_p norms, convex distance functions) or different kinds of sites (points, segments, circles) could all be handled using similar techniques. To formalize this, several abstract frameworks were defined, such as the one of Edelsbrunner and Seidel [9] and the two variants of abstract Voronoi diagrams of Klein [10, 12]. In this paper we define a new abstract framework to deal with Voronoi diagrams constructed incrementally by inserting new sites.

Let G be a 3-regular embedded plane graph with n vertices[1]. We seek to bound the number of edge removals and insertions needed to implement the following operation, hereafter referred to as a *flarb*[2]: Given a simple closed curve $\mathcal{C}$ in the plane whose interior intersects G in a connected component, split both $\mathcal{C}$ and all the edges that it crosses at the point of intersection, remove every edge and vertex that lies in the interior of $\mathcal{C}$, and add each curve in the subdivision of $\mathcal{C}$ as a new edge; see Figure 1. This operation can be used to represent the insertion of new cells in different types of Voronoi diagrams. It can also be used to represent the changes to the 1-skeleton of a polyhedron in $\mathbb{R}^3$ after it is intersected with a halfspace.

Results

We show that the amortized *cost* of a flarb operation, where the combinatorial cost is defined to be the minimum number of edge insertions and removals needed to perform it, is $O(\sqrt{n})$. We also show a matching lower bound: some sequences of flarbs require $\Omega(\sqrt{n})$ links and cuts per flarb, even when the graph is a tree (or more precisely a Halin graph—a tree with all leaves connected by a cycle to make it 3-regular). This contrasts with the $O(\log n)$ upper bound of Aronov et al. [2] for the Voronoi diagram of points in convex position (also a tree) when points are added in clockwise order.

We complement these combinatorial bounds with an algorithmic result. We present an output-sensitive data structure that maintains the nearest- or farthest-point Voronoi diagram of a set S of n points in convex

[1] While the introduction used n for the number of sites in S, the combinatorial part of this article uses n for the number of vertices in the Voronoi diagram. By Euler's formula, those two values are asymptotically equivalent, up to a constant factor.

[2] Although the last two authors are honored by the flattering renaming of the flarb operation in the literature [14], this paper uses original terminology.

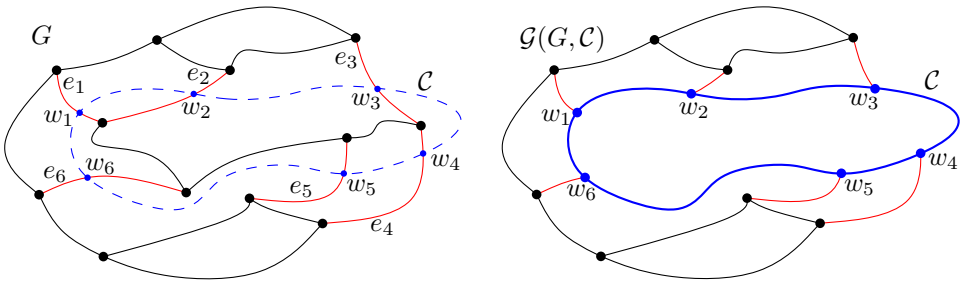

Figure 1: The flarb operation on a graph G induced by a flarbable curve $\mathcal{C}$, produces a graph $\mathcal{G}(G,\mathcal{C})$ with 2 more vertices. Fleeq-edges crossed by $\mathcal{C}$ are shown in red.

position as new points are added to S. Upon an insertion, the data structure finds the minimum (up to within a constant factor) number K of edge insertions and deletions necessary to update the Voronoi diagram of S.

The running time of each insertion is $O(K \log^7 n)$, and by our combinatorial bounds, $K = O(\sqrt{n})$. This solves the open problem posed by Aronov et al. [2].

The distinguishing feature of this data structure is that it explicitly maintains the graph structure of the Voronoi diagram after every insertion, a property that is not provided by any nearest neighbor data structure that uses decomposable searching problem techniques. Further, the data structure also maintains the Voronoi diagram in a *grappa tree* [2], a variant of the link-cut tree of Sleator and Tarjan [15] which allows a powerful query operation called *oracle-search*. Roughly speaking, the oracle-search query has access to an oracle specifying a vertex to find. Given an edge of the tree, the oracle determines which of the two subtrees attached to its endpoints contains that vertex. Grappa trees use $O(\log n)$ time and oracle calls to find the sought vertex. A grappa tree is in some sense a dynamic version of the centroid decomposition for trees, which is used in many algorithms for searching in Voronoi diagrams. Using this structure, it is possible to solve a number of problems for the set S at any moment during the incremental construction, for example:

- Given sites p and q, report whether they are connected by a Delaunay edge in $O(\log n)$ time.
- Given a point q, find the Voronoi cell containing q in $O(\log n)$ time. This not only gives the nearest neighbor of q, but a pointer to the explicit description of its cell.
- Find the smallest disk enclosing S, centered on a query segment $[pq]$, in $O(\log n)$ time [5].
- Find the smallest disk enclosing S, centered on a query circle C, in $O(\log n)$ time [3].
- Given a convex polygon P (counterclockwise array of its m vertices), find the smallest disk enclosing S and excluding P in $O(\log n + \log m)$ time [1].

The combinatorial bound for Voronoi diagrams also has direct algorithmic consequences, the most important being that it is possible to store all versions of the graph throughout a sequence of insertions using persistence in $O(n^{3/2})$ space. Since the entire structure of the graph is stored for each version, this provides a foundation for many applications that, for instance would require searching the sequence of insertions for the moment during which a specific event occurred.

Outline

The main approach used to bound the combinatorial cost of a flarb is to examine how the complexity of the faces changes. Notice that faces whose size remains the same do not require edge insertions and deletions. The other faces either grow or shrink, and a careful counting argument reveals that the cost of a flarb is at most the number faces that shrink (or disappear) upon execution of the flarb (Section 2). By using a potential function that sums the sizes of all faces, the combinatorial cost of shrinking faces is paid for by the reduction of their potential. To avoid incurring a high increase in potential for a large new face, the potential

of each face is capped at $\sqrt{n}$. Then at most $O(\sqrt{n})$ large faces can shrink without changing potential and are accounted for separately (Section 3). The matching $\Omega(\sqrt{n})$ lower bound is presented in Section 4, and Section 5 presents the data structure for performing flarbs for the Voronoi diagrams of points in convex position.

2 The flarb operation

In this section we formalize the flarb operation that models the insertion of new sites in Voronoi diagrams and present a preliminary analysis of the cost of a flarb.

Let $G = (V, E)$ be a planar 3-regular graph embedded in $\mathbb{R}^2$ (not-necessarily with a straight-line embedding). Let $\mathcal{C}$ be a simple closed Jordan curve in the plane. Define $\textsc{in}(\mathcal{C})$ to be the set of vertices of G that lie in the interior of $\mathcal{C}$ and let $\textsc{ex}(\mathcal{C}) = V \setminus \textsc{in}(\mathcal{C})$. We say that $\mathcal{C}$ is *flarbable* for G if the following conditions hold:

1. the graph induced by $\textsc{in}(\mathcal{C})$ is connected,
2. $\mathcal{C}$ intersects each edge of G either at a single point or not at all,
3. $\mathcal{C}$ passes through no vertex of G, and
4. the intersection of $\mathcal{C}$ with each face of G is path-connected.

In the case where the graph G is clear from context, we simply say that $\mathcal{C}$ is flarbable. The *fleeq* of $\mathcal{C}$ is the circular sequence $\mathcal{E}_\mathcal{C} = e_1, \ldots, e_k$ of edges in E that are crossed by $\mathcal{C}$; we call the edges in $\mathcal{E}_\mathcal{C}$ *fleeq-edges*. A face whose interior is crossed by $\mathcal{C}$ is called a *$\mathcal{C}$-face*. We assume without loss of generality that $\mathcal{C}$ is oriented clockwise and that the edges in $\mathcal{E}_\mathcal{C}$ are ordered according to their intersection with $\mathcal{C}$. Given a flarbable curve $\mathcal{C}$ on G, we present the following definition.

Definition 2.1. For a planar graph G and a curve $\mathcal{C}$ that is flarbable for G, we define a *flarb* operation $\mathcal{F}(G, \mathcal{E}_\mathcal{C})$ which produces a new 3-connected graph $\mathcal{G}(G, \mathcal{C})$ as follows (see Figure 1 for a depiction):

1. For each edge $e_i = (u_i, v_i)$ in $\mathcal{E}_\mathcal{C}$ such that $u_i \in \textsc{in}(\mathcal{C})$ and $v_i \in \textsc{ex}(\mathcal{C})$, create a new vertex $w_i = \mathcal{C} \cap e_i$ and connect it to v_i along e_i.
2. For each pair e_i, e_{i+1} of successive edges in $\mathcal{E}_\mathcal{C}$, create a new edge (w_i, w_{i+1}) between them along $\mathcal{C}$. We call (w_i, w_{i+1}) a *$\mathcal{C}$-edge* (all indices are taken modulo k).
3. Delete all vertices of $\textsc{in}(\mathcal{C})$ along with their incident edges.

Lemma 2.2. *For each flarbable curve $\mathcal{C}$ on a 3-regular planar graph G, $\mathcal{G}(G, \mathcal{C})$ has at most 2 more vertices than G does.*

Proof. Let $\mathcal{E}_\mathcal{C} = e_1, \ldots, e_k$ be the fleeq of $\mathcal{C}$ and let f be the new face in $\mathcal{G}(G, \mathcal{C})$ that is bounded by $\mathcal{C}$ and created by the flarb operation $\mathcal{F}(G, \mathcal{E}_\mathcal{C})$. Notice that the vertices of f are the points $w_1, \ldots, w_k$ along edges $e_1, \ldots, e_k$, where $w_i = \mathcal{C} \cap e_i$. Since $\mathcal{C}$ is flarbable, the subgraph induced by the vertices of $\textsc{in}(\mathcal{C}) \cup \{w_1, \ldots, w_k\}$ is also a connected graph T with $w_1, \ldots, w_k$ as its leaves and every other vertex of degree 3; see Figure 1. Therefore T has at least $k - 2$ internal vertices. The flarb operation adds k vertices, namely $w_1, \ldots, w_k$, and the internal vertices of T are deleted. Therefore, the net increase in the number of vertices is at most 2. □

Since each newly created vertex has degree three and all remaining vertices are unaffected, the new graph is 3-regular. In other words, the flarb operation $\mathcal{F}(G, \mathcal{E}_\mathcal{C})$ creates a cycle along $\mathcal{C}$ and removes the portion of the graph enclosed by $\mathcal{C}$. Note that for any point set in general position (no four points lie on the same circle), its Voronoi diagram is a 3-regular planar graph, assuming we use the line at infinity to join the endpoints of its unbounded edges in clockwise order. Therefore, a flarb can be used to represent the changes to the Voronoi diagram upon insertion of a new site.

Observation 2.3. *Given a set S of points in general position, let $\mathcal{V}(S)$ be the graph of the Voronoi diagram of S. For a new point q, there exists some curve $\mathcal{C}_S^q$ such that $\mathcal{G}(\mathcal{V}(S), \mathcal{C}_S^q) = \mathcal{V}(S \cup \{q\})$; namely, $\mathcal{C}_S^q$ is the boundary of the Voronoi cell of q in $\mathcal{V}(S \cup \{q\})$.*

More generally, convex polytopes defined by the intersection of halfspaces in $\mathbb{R}^3$ behave similarly: the intersection of a new halfspace with a convex polytope modifies the structure of its 1-skeleton by adding a new face. This structural change can be implemented by performing a flarb operation in which the flarbable curve consists of the boundary of the new face.

Preserved faces and edges

Definition 2.4. Given a $\mathcal{C}$-face f of G, the *modified face* of f is the face f' of $\mathcal{G}(G, \mathcal{C})$ that coincides with f outside of $\mathcal{C}$. In other words, f' is the face that remains from f after performing the flarb $\mathcal{F}(G, \mathcal{E}_\mathcal{C})$. We say that a $\mathcal{C}$-face f is *preserved* (by the flarb $\mathcal{F}(G, \mathcal{E}_\mathcal{C})$) if $|f| = |f'|$. Moreover, we say that each edge in a preserved face is *preserved* (by $\mathcal{F}(G, \mathcal{E}_\mathcal{C})$). Denote by $\mathcal{P}(G, \mathcal{C})$ the set of faces preserved by $\mathcal{F}(G, \mathcal{E}_\mathcal{C})$ and let $\mathcal{B}(G, \mathcal{C})$ be the set of faces wholly contained in the interior of $\mathcal{C}$.

Since a preserved $\mathcal{C}$-face bounded by two fleeq-edges e_i and e_{i+1} has the same size before and after the flarb, there must be an edge e of G connecting e_i with e_{i+1} which is replaced by a $\mathcal{C}$-edge e^* after the flarb. In this case, we say that the edge e *reappears* as e^*.

The following auxiliary lemma will help us bound the number of operations needed to produce the graph $\mathcal{G}(G, \mathcal{C})$, and follows directly from the Euler characteristic of connected planar graphs:

Lemma 2.5. *Let H be a connected planar graph with vertices of degree either 1, 2 or 3. For each $i \in \{1, 2, 3\}$, let δ_i be the number of vertices of H with degree i. Then, H has exactly $2\delta_1 + \delta_2 + 3F_H - 3$ edges, where F_H is the number of bounded faces of H.*

2.1 Combinatorial cost of a flarb

Given a 3-regular graph $G = (V, E)$ and a flarbable curve $\mathcal{C}$ we want to analyze the number of structural changes that G must undergo to perform $\mathcal{F}(G, \mathcal{E}_\mathcal{C})$. To this end, we define the combinatorial *cost* of $\mathcal{F}(G, \mathcal{E}_\mathcal{C})$, denoted by $\textsc{cost}(G, \mathcal{C})$, to be the minimum number of links and cuts needed to transform G into $\mathcal{G}(G, \mathcal{C})$ (note that the algorithm may not implement the flarb operation according to the procedure described in Definition 2.1). We assume that any other operation has no cost and is therefore not included in the cost of the flarb.

Consider the fleeq $\mathcal{E}_\mathcal{C} = e_1, \ldots, e_k$ and the $\mathcal{C}$-edges created by $\mathcal{F}(G, \mathcal{E}_\mathcal{C})$. Let e be an edge adjacent to some e_i and e_{i+1} that reappears as the $\mathcal{C}$-edge e^*. Notice that we can obtain e^* without any links or cuts to G: simply shrink e_i and e_{i+1} so that their endpoints in $\textsc{in}(\mathcal{C})$ now coincide with their intersections with $\mathcal{C}$. Then modify e to coincide with the portion of $\mathcal{C}$ connecting the new endpoints of e_i and e_{i+1}. Using this *preserving operation*, we obtain the $\mathcal{C}$-edge e^* with no cost to the flarb. Intuitively, preserved edges are cost-free in a flarb while non-preserved edges have a nonzero cost. This notion is formalized in the following lemma.

Lemma 2.6. *For a flarbable curve $\mathcal{C}$,*

$$(|\mathcal{E}_\mathcal{C}| + |\mathcal{B}(G, \mathcal{C})| - |\mathcal{P}(G, \mathcal{C})|)/2 \leq \textsc{cost}(G, \mathcal{C}) \leq 4|\mathcal{E}_\mathcal{C}| + 3|\mathcal{B}(G, \mathcal{C})| - 4|\mathcal{P}(G, \mathcal{C})|.$$

Proof. For the upper bound, we describe a construction of $\mathcal{G}(G, \mathcal{C})$ from G using at most $|\mathcal{E}_\mathcal{C}| + 3|\mathcal{B}(G, \mathcal{C})| - 4|\mathcal{P}(G, \mathcal{C})|$ links and cuts[3]. Consider the subgraph $\mathcal{G}_\mathcal{C}$ induced by $\textsc{in}(\mathcal{C}) \cup \{v : v \text{ is an endpoint of some edge in } \mathcal{E}_\mathcal{C}\}$. Since $\mathcal{C}$ is flarbable, $\mathcal{G}_\mathcal{C}$ is a connected graph such that each vertex of $\textsc{in}(\mathcal{C})$ has degree 3 while the endpoints of the fleeq-edges outside of $\mathcal{C}$ have degree 1. Note that if two preserved faces share a non-fleeq edge e, then there are four neighbors of the endpoints of e that lie outside of $\mathcal{C}$. Since $\mathcal{G}_\mathcal{C}$ is connected, e and its four adjacent edges define the entire graph $\mathcal{G}_\mathcal{C}$ and the bound holds trivially. Therefore, we assume that no two

[3]We caution the reader that while this construction is algorithmic in nature, it is used purely to provide an upper-bound and does not reflect the behavior of the algorithm presented in Section 5 that gives our desired runtime.

preserved faces share a non-fleeq-edge from this point forward.

Note that the bounded faces of $\mathcal{G}_\mathcal{C}$ are exactly the bounded faces in $\mathcal{B}(G,\mathcal{C})$. Since $\mathcal{G}_\mathcal{C}$ has $|\mathcal{E}_\mathcal{C}|$ vertices of degree 1, no vertices of degree 2, and $|\mathcal{B}(G,\mathcal{C})|$ bounded faces, by Lemma 2.5, $\mathcal{G}_\mathcal{C}$ has at most $2|\mathcal{E}_\mathcal{C}|+3|\mathcal{B}(G,\mathcal{C})|$ edges. Every edge of $\mathcal{G}_\mathcal{C}$ that is not preserved is removed with a cut operation (isolated vertices will be removed afterwards). Note that each preserved face contains at least three preserved edges: two fleeq-edges and a third edge of G. Based on the assumption that no two preserved faces share a non-fleeq-edge, the third edge is not double counted, while the fleeq-edges may be counted at most twice. Therefore, each preserved face contributes at least two preserved edges that are specific to that face, meaning that a total of at most $2|\mathcal{E}_\mathcal{C}| + 3|\mathcal{B}(G,\mathcal{C})| - 2|\mathcal{P}(G,\mathcal{C})|$ cut operations are performed. Note that each non-preserved fleeq-edge has been cut and will need to be reintroduced later to obtain $\mathcal{G}(G,\mathcal{C})$.

Recall that no edge bounding a preserved face has been cut. For each preserved face, perform a preserving operation on it which requires no link or cut operation. Since no two preserved faces share a non-fleeq edge, all the $\mathcal{C}$-edges bounding the preserved faces are added without increasing $\textsc{cost}(G,\mathcal{C})$. To complete the construction of $\mathcal{G}(G,\mathcal{C})$, create each fleeq-edge that is not preserved and then add the remaining $\mathcal{C}$-edges bounding non-preserved $\mathcal{C}$-faces. Because at least $|\mathcal{P}(G,\mathcal{C})|$ fleeq-edges were preserved, at most $|\mathcal{E}_\mathcal{C}|-|\mathcal{P}(G,\mathcal{C})|$ fleeq-edges must be reintroduced. Moreover, since only $|\mathcal{E}_\mathcal{C}| - |\mathcal{P}(G,\mathcal{C})|$ $\mathcal{C}$-faces are not preserved, we need to create at most $|\mathcal{E}_\mathcal{C}| - |\mathcal{P}(G,\mathcal{C})|$ $\mathcal{C}$-edges. Therefore, this last step completes the flarb and construct $\mathcal{G}(G,\mathcal{C})$ using a total of at most $2|\mathcal{E}_\mathcal{C}| - 2|\mathcal{P}(G,\mathcal{C})|$ link operations. Consequently, the total number of link and cuts needed to obtain $\mathcal{G}(G,\mathcal{C})$ from G is at most $4|\mathcal{E}_\mathcal{C}| + 3|\mathcal{B}(G,\mathcal{C})| - 4|\mathcal{P}(G,\mathcal{C})|$ as claimed.

To show that $\textsc{cost}(G,\mathcal{C}) > (|\mathcal{E}_\mathcal{C}| + |\mathcal{B}(G,\mathcal{C})| - |\mathcal{P}(G,\mathcal{C})|)/2$, simply note that in every non-preserved $\mathcal{C}$-face, the algorithm needs to perform at least one cut, either to augment the size or reduce the size of the face. Because G and $\mathcal{C}$ define exactly $|\mathcal{E}_\mathcal{C}| + |\mathcal{B}(G,\mathcal{C})|$ faces, and since in all but $|\mathcal{P}(G,\mathcal{C})|$ of them at least one of its edges must be cut, at least $|\mathcal{E}_\mathcal{C}| + |\mathcal{B}(G,\mathcal{C})| - |\mathcal{P}(G,\mathcal{C})|$ edges must be cut. Since an edge belongs to at most two faces a cut can be over-counted at most twice and the claimed bound holds. □

3 The combinatorial upper bound

In this section, we define a potential function to bound the amortized cost of each operation in a sequence of flarb operations. For a 3-regular embedded planar graph $G = (V, E)$, we define two potential functions: a local potential function μ to measure the potential of each face, and a global potential function Φ to measure the potential of the whole graph.

Definition 3.1. Let F be the set of faces of a 3-regular embedded planar graph $G = (V, E)$. For each face $f \in F$, let $\mu(f) = \min\{\lceil\sqrt{|V|}\rceil, |f|\}$, where $|f|$ is the number of edges on the boundary of f. The potential $\Phi(G)$ of G is defined as follows:

$$\Phi(G) = \lambda \sum_{f\in F} \mu(f),$$

for some sufficiently large positive constant λ to be defined later.

Recall that the potential $\mu(f)$ of a $\mathcal{C}$-face f remains unchanged as long as $|f|, |f'| \geq \sqrt{|V|}$, where f' is the modified face of f after the flarb. Since there is no change in potential that we can use within large $\mathcal{C}$-faces, we exclude them from our analysis and focus only on smaller $\mathcal{C}$-faces. We formalize this notion in the following section.

3.1 Flarbable sub-curves

Given a flarbable curve $\mathcal{C}$, a (connected) curve $\gamma \subseteq \mathcal{C}$ is a *flarbable sub-curves*. Let $\epsilon_\gamma = e_1, \ldots, e_k$ (or simply ϵ) be the set of fleeq-edges intersected by γ given in order of intersection after orienting γ arbitrarily. We call ϵ the *subfleeq* induced by γ. We say that a face is a γ-face if two of its fleeq-edges are crossed by γ (if γ has an endpoint in the interior of this face, it is not a γ-face).

Consider the set of all edges of G intersected or enclosed by $\mathcal{C}$ that bound some γ-face. Since $\mathcal{E}_\mathcal{C}$ is

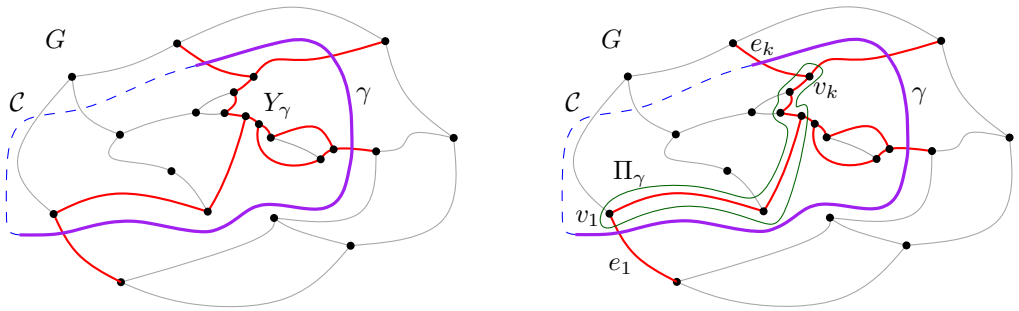

Figure 2: Left: A flarbable sub-curves γ is contained in a flarbable curve $\mathcal{C}$. The graph Y_γ is the union of all edges bounding a γ-face. Right: The path Π_γ connects the endpoints of the first and last fleeq-edges crossed by γ by going along the boundary of the outer-face of Y_γ.

flarbable, these edges induce a connected subgraph Y_γ of G with $|\epsilon| = k$ leaves (vertices of degree 1), namely the endpoints outside of $\mathcal{C}$ of each fleeq-edge in ϵ; see Figure 2. Notice that Y_γ may consist of some bounded faces contained in the interior of $\mathcal{C}$. Let H_γ be the set of bounded faces of Y_γ and let δ_2 be the number of vertices of degree 2 of Y_γ. Since Y_γ consists of k vertices of degree 1, Lemma 2.5 implies the following result.

Corollary 3.2. *The graph Y_γ consists of exactly $2k + \delta_2 + 3|H_\gamma| - 3$ edges.*

Recall that a $\mathcal{C}$-face f is preserved if its corresponding modified face f' in $\mathcal{G}(G, \mathcal{C})$ has the same number of edges, i.e., $|f'| = |f|$. We say that f is *augmented* if $|f'| = |f| + 1$ and we call f *shrinking* if $|f'| < |f|$. Notice that these are all the possible cases as f gains at most one new edge during the flarb, namely the $\mathcal{C}$-edge crossing this face.

In the context of a particular flarbable sub-curve γ, let a_γ, s_γ and p_γ be the number of augmented, shrinking and preserved γ-faces, respectively (or simply a, s and p if γ is clear from the context). We further differentiate among the s shrinking γ-faces. A shrinking γ-face is *interior* if it contains no vertex of degree 2 of Y_γ and does not share an edge with an augmenting face. Let s_a be the number of shrinking γ-faces that share an edge with an augmented face, let s_b be the number of shrinking γ-faces not adjacent to an augmented face that have a vertex of degree 2 of Y_γ, and let s_c be the number of interior shrinking γ-faces. Therefore, $s = s_a + s_b + s_c$ is the total number of shrinking γ-faces.

Since each augmented face has at most two edges and because there are a augmented faces, we know that $s_a \leq 2a$. Let v_1 and v_k be the endpoints of the edges e_1 and e_k that lie inside $\mathcal{C}$. Let Π_γ be the unique path connecting v_1 and v_k in Y_γ that traverses the boundary of the outer face of Y_γ and stays in the interior of $\mathcal{C}$; see Figure 2.

Notice that Π_γ contains all the edges of γ-faces that may bound a γ'-face for some other flarbable sub-curve γ' disjoint from γ. In the end, we aim to have bounds on the number of edges that will be removed from the γ-faces during the flarb, but some of these edges may be double counted if they are shared with a γ'-face. Therefore, we aim to bound the length of Π_γ and count precisely the number of edges that could possibly be double-counted.

Lemma 3.3. *The path Π_γ has length at most $k + 3|H_\gamma| + \delta_2 - a - s_c$.*

Proof. Notice that no fleeq-edge can be part of Π_γ or this path would go outside of $\mathcal{C}$, i.e., there are k fleeq-edges of Y_γ that cannot be part of Π_γ.

We say that a vertex is *augmented* if it is incident to two fleeq-edges and a third edge that is not part of ϵ, which we call an *augmented edge*. Because each augmented γ-face has exactly one augmented vertex, there are exactly a augmented vertices in Y_γ. Moreover, Π_γ contains at most 2 augmented vertices (if v_1 or v_k is augmented). Thus, at most two augmented edges can be traversed by Π_γ and hence, at least $a - 2$ augmented edges of Y_γ do not belong to Π_γ.

Let f be an internal shrinking γ-face. Since f is not adjacent to an augmented γ-face, it has no augmented

edge on its boundary. We claim that f has at least one edge that is not traversed by Π_γ. If this claim is true, then there are at least s_c non-fleeq non-augmented edges that cannot be used by Π_γ—one for each internal shrinking γ-face. Thus, since Y_γ consists of $2k + 3|H_\gamma| + \delta_2 - 2$ edges, the number of edges in Π_γ is at most

$$2k + 3|H_\gamma| + \delta_2 - 2 - (k + a - 2 + s_c) = k + 3|H_\gamma| + \delta_2 - a - s_c.$$

It remains to show that each internal shrinking γ-face f has at least one non-fleeq edge that is not traversed by Π_γ. If Π_γ contains no edge on the boundary of f, then the claim holds trivially. If Π_γ contains exactly one edge of f, then since f is shrinking, it has at least 4 edges and two of them are not fleeq-edges. Thus, in this case there is one edge of f that is not traversed by Π_γ. We assume from now on that Π_γ contains at least two edges of f.

We claim that that Π_γ visits a contiguous sequence of edges along the boundary of f. To see this, note that each face of Y_γ lying between Π_γ and the boundary of f cannot be crossed by $\mathcal{C}$. Therefore, if we consider the first edge of Π_γ that is not on the boundary of f after visiting f the first time, the this edge is incident to the outer face of Y_γ and the only face of Y_γ that it is incident with does not intersect $\mathcal{C}$. This is a contradiction, since this edge should not be part of Y_γ by definition. Therefore, Π_γ visits a contiguous sequence of edges along f.

If Π_γ visits 2 consecutive edges of f, then the vertex in between them must have degree 2 in Y_γ, as the two edges are incident to the outer face—a contradiction since f is an internal shrinking face with no vertex of degree 2. Consequently, if f is an internal shrinking face, it has always at least one non-fleeq edge that is not traversed by Π_γ. □

3.2 How much do faces shrink in a flarb?

In order to analyze the effect of the flarb operations on flarbable sub-curves, we think of each edge as consisting of two *half-edges*, each adjacent to one of the two faces incident to this edge. For a given edge, the algorithm may delete its half-edges during two separate flarbs of differing flarbable sub-curves.

We define the operation $\mathcal{F}(G, \gamma)$ to be the operation which executes steps 1 and 2 of the flarb on the flarbable sub-curve γ and then deletes each half-edge with both endpoints in $\text{IN}(\mathcal{C})$ adjacent to a γ-face. Since $\mathcal{F}(G, \gamma)$ removes and adds half-edges, we are interested in bounding the net balance of half-edges throughout the flarb. To do this, we measure the change in size of a face during the flarb.

Recall that a, s and p are the number of augmented, shrinking and preserved γ-faces, respectively. The following result provides a bound on the total "shrinkage" of the faces crossed by a given flarbable sub-curve.

Theorem 3.4. *Given a flarbable curve $\mathcal{C}$ on G and a flarbable sub-curve γ crossing the fleeq-edges $\epsilon = e_1, \ldots, e_k$, let $f_1, \ldots, f_k$ be the sequence of γ-faces and let $f'_1, \ldots, f'_k$ be their corresponding modified faces after the flarb $\mathcal{F}(G, \gamma)$. Then,*

$$\sum_{i=1}^{k} (|f_i| - |f'_i|) \geq s/2. \tag{1}$$

Proof. Recall that no successive pair of γ-faces can both be augmented unless $\mathcal{E}_\mathcal{C}$ consists of three edges incident to a single vertex. In this case, at most 3 γ-faces can be augmented, so $\sum_{i=1}^{k} (|f_i| - |f'_i|) = 3$ and the result holds trivially; hence, we assume from now on that no two successive faces are both augmented.

Let Δ be the number of half-edges removed during $\mathcal{F}(G, \gamma)$. Notice that to count how much a face f_i shrinks when becoming f'_i after the flarb, we need to count the number of half-edges of f_i that are deleted and the number that are added in f'_i. Since exactly one half-edge is added in each f'_i, we know that $\sum_{i=1}^{k} (|f_i| - |f'_i|) = \Delta - k$. We claim that $\Delta \geq k + s/2$. If this claim is true, then $\sum_{i=1}^{k} (|f_i| - |f'_i|) \geq s/2$ as stated in the theorem. In the remainder of this proof, we show this bound on Δ.

Let $\mathcal{T} = (V_\mathcal{T}, E_\mathcal{T})$ be the subgraph of Y_γ obtained by removing its k fleeq-edges. It follows from 3.2 that $|E_\mathcal{T}| = k + 3|H_\gamma| + \delta_2 - 3$. To obtain a precise counting of Δ, notice that for some edges of $\mathcal{T}$, $\mathcal{F}(G, \gamma)$ removes only one of their half-edges and for others it will remove both of them. Since the fleeq-edges are present in each of the faces $f_1, \ldots, f_k$ before and after the flarb, we get that

$$\Delta = 2|E_\mathcal{T}| - S_\mathcal{T}, \tag{2}$$

where $S_{\mathcal{T}}$ denotes the number of edges in $\mathcal{T}$ with only one half-edge incident to a face of $f_1, \ldots, f_k$.

Note that the edges of $S_{\mathcal{T}}$ are exactly the edges on the path Π_γ bounded in Lemma 3.3. Therefore, $S_{\mathcal{T}} \leq k + 3|H_\gamma| + \delta_2 - a - s_c$. By using this bound in (2), we get

$$\Delta \geq 2(k + 3|H_\gamma| + \delta_2 - 3) - (k + 3|H_\gamma| + \delta_2 - a - s_c) = k + 3|H_\gamma| + \delta_2 + a + s_c - 6.$$

Since each shrinking γ-face accounted for by s_b has a vertex of degree 2 in Y_γ, we know that $\delta_2 \geq s_b$. Moreover, $s_a \leq 2a$ as each shrinking γ-face can be adjacent to at most two augmenting γ-faces. Therefore, since $s = s_a + s_b + s_c$, we get that $\Delta \geq k + 3|H_\gamma| + s_a/2 + s_b + s_c \geq k + s/2$, where s is the number of shrinking γ-faces proving the claimed bound on Δ. □

3.3 Flarbable sequences

Let $\mathcal{G}^0 = G$. A sequence of curves $\mathscr{C} = \mathcal{C}_1, \ldots, \mathcal{C}_k$ is *flarbable* if for each $i \in [k]$, $\mathcal{C}_i$ is a flarbable on

$$\mathcal{G}^i = \mathcal{G}(\mathcal{G}^{i-1}, \mathcal{C}_i).$$

As a notational shorthand, let $\mathcal{F}^i$ denote the flarb operation $\mathcal{F}(\mathcal{G}^{i-1}, \mathcal{C}_i)$ when $\mathscr{C}$ is a flarbable sequence for G.

Theorem 3.5. *For a 3-regular planar graph $G = (V, E)$ and some flarbable sequence $\mathscr{C} = \mathcal{C}_1, \ldots, \mathcal{C}_N$ of flarbable fleeqs, for all $i \in [N]$,*

$$\textsc{cost}(\mathcal{G}^{i-1}, \mathcal{C}_i) + \Phi(\mathcal{G}^i) - \Phi(\mathcal{G}^{i-1}) \leq O(\sqrt{|V_i|}),$$

where V_i is the set of vertices of $\mathcal{G}^i$.

Proof. Partition $\mathcal{C}_i$ into smaller curves $\gamma_1, \ldots, \gamma_h$ such that for all $j \in [h]$, γ_j is a maximal curve contained in $\mathcal{C}_i$ that does not intersect the interior of a face with more than $\sqrt{|V_i|}$ edges. Since there can be at most $\sqrt{|V_i|}$ faces of size $\sqrt{|V_i|}$, we know that $h \leq \sqrt{|V_i|}$. Let ϵ_j be the subfleeq containing each fleeq-edge crossed by γ_j. Let a_j, s_j and p_j be the number of augmented, shrinking and preserved γ_j-faces, respectively. Notice that $|\epsilon_j| = a_j + s_j + p_j + 1$. Moreover, since each augmented face is adjacent to a shrinking face, we know that $a_j \leq s_j + 1$. Therefore, $|\epsilon_j| \leq 2s_j + p_j + 2$.

Let $\mathcal{L}_i$ be the set of $\mathcal{C}_i$-faces with at least $\sqrt{|V_i|}$ edges and let ω_i be the set of all faces of $\mathcal{G}^{i-1}$ completely enclosed in the interior of $\mathcal{C}_i$.

First, we upper bound $\textsc{cost}(\mathcal{G}^{i-1}, \mathcal{C}_i)$. By Lemma 2.6, we know that

$$\begin{aligned}
\textsc{cost}(\mathcal{G}^{i-1}, \mathcal{C}_i) &\leq 4|\mathcal{E}_{\mathcal{C}_i}| + 3|\mathcal{B}(\mathcal{G}^{i-1}, \mathcal{C}_i)| - 4|\mathcal{P}(\mathcal{G}^{i-1}, \mathcal{C}_i)| && (3)\\
&= 4\sum_{j=1}^{h} |\epsilon_j| + 3|\mathcal{B}(\mathcal{G}^{i-1}, \mathcal{C}_i)| - 4|\mathcal{P}(\mathcal{G}^{i-1}, \mathcal{C}_i)| && (4)\\
&\leq 4\sum_{j=1}^{h} (2s_j + p_j + 2) + 3|\mathcal{B}(\mathcal{G}^{i-1}, \mathcal{C}_i)| - 4|\mathcal{P}(\mathcal{G}^{i-1}, \mathcal{C}_i)| && (5)
\end{aligned}$$

because each preserved face is crossed by exactly one flarbable sub-curve, $\sum_{j=1}^{h} p_j = |\mathcal{P}(\mathcal{G}^{i-1}, \mathcal{C}_i)|$. Therefore,

$$\textsc{cost}(\mathcal{G}^{i-1}, \mathcal{C}_i) \leq 4\sum_{j=1}^{h} (2s_j + 2) + 3|\mathcal{B}(\mathcal{G}^{i-1}, \mathcal{C}_i)| = 8h + 8\sum_{j=1}^{h} s_j + 3|\mathcal{B}(\mathcal{G}^{i-1}, \mathcal{C}_i)| \,.$$

Since $h \leq \sqrt{|V_i|}$, we conclude that

$$\textsc{cost}(\mathcal{G}^{i-1}, \mathcal{C}_i) \leq 8\sqrt{|V_i|} + 8\sum_{j=1}^{h} s_j + 3|\mathcal{B}(\mathcal{G}^{i-1}, \mathcal{C}_i)| \,. \quad (6)$$

Next, we upper bound the change in potential $\Phi(\mathcal{G}^i) - \Phi(\mathcal{G}^{i-1})$. Given a flarbable curve or sub-curve γ, let $\mathcal{A}(\gamma)$ denote the set of γ-faces. Recall that for a γ-face $f \in \mathcal{A}(\gamma)$, f' is the modified face of f. Also, let f_n be the new face created by $\mathcal{F}^i$, i.e., the face of $\mathcal{G}^i$ bounded by $\mathcal{C}_i$. Recall that for each face $f \in \mathcal{B}(\mathcal{G}^{i-1}, \mathcal{C}_i)$, f is removed and the potential decreases by $\mu(f) \geq 1$. Using this, we can break up the summation to obtain the following:

$$\Phi(\mathcal{G}^i) - \Phi(\mathcal{G}^{i-1}) = \mu(f_n) + \lambda \sum_{f \in \mathcal{A}(\mathcal{C}_i)} (\mu(f') - \mu(f)) - \lambda \sum_{f \in \mathcal{B}(\mathcal{G}^{i-1}, \mathcal{C}_i)} \mu(f) \tag{7}$$

$$\leq \mu(f_n) + \lambda \sum_{f \in \mathcal{A}(\mathcal{C}_i)} (\mu(f') - \mu(f)) - \lambda |\mathcal{B}(\mathcal{G}^{i-1}, \mathcal{C}_i)| \ . \tag{8}$$

We now break up the first summation by independently considering the large faces in $\mathcal{L}_i$ and the remaining smaller faces which are crossed by some flarbable sub-curve. Then

$$\Phi(\mathcal{G}^i) - \Phi(\mathcal{G}^{i-1}) \leq \mu(f_n) + \lambda \sum_{j=1}^{h} \left(\sum_{f \in \mathcal{A}(\gamma_j)} (\mu(f') - \mu(f)) \right) \tag{9}$$

$$+ \lambda \sum_{f \in \mathcal{L}_i} (\mu(f') - \mu(f)) - \lambda |\mathcal{B}(\mathcal{G}^{i-1}, \mathcal{C}_i)|. \tag{10}$$

Since each face can gain at most one edge, in particular we know that $\mu(f') - \mu(f) \leq 1$ for each $f \in \mathcal{L}_i$. Moreover, $\mu(f_n) \leq \sqrt{|V_i|}$ by definition. Thus,

$$\Phi(\mathcal{G}^i) - \Phi(\mathcal{G}^{i-1}) \leq \sqrt{|V_i|} + \lambda \sum_{j=1}^{h} \left(\sum_{f \in \mathcal{A}(\gamma_j)} (\mu(f') - \mu(f)) \right) + \lambda |\mathcal{L}_i| - \lambda |\mathcal{B}(\mathcal{G}^{i-1}, \mathcal{C}_i)|.$$

Note that $\mu(f) = |f|$ for each face $f \in \mathcal{A}(\gamma_j)$, $1 \leq j \leq h$. Thus, applying Theorem 3.4 to the first summation, we get

$$\Phi(\mathcal{G}^i) - \Phi(\mathcal{G}^{i-1}) \leq \sqrt{|V_i|} - \frac{\lambda}{2} \sum_{j=1}^{h} s_j + \lambda |\mathcal{L}_i| - \lambda |\mathcal{B}(\mathcal{G}^{i-1}, \mathcal{C}_i)|.$$

Since there can be at most $\sqrt{|V_i|}$ faces of size $\sqrt{|V_i|}$, we know that $|\mathcal{L}_i| \leq \sqrt{|V_i|}$. Therefore,

$$\Phi(\mathcal{G}^i) - \Phi(\mathcal{G}^{i-1}) \leq (\lambda + 1)\sqrt{|V_i|} - \frac{\lambda}{2} \sum_{j=1}^{h} s_j - \lambda |\mathcal{B}(\mathcal{G}^{i-1}, \mathcal{C}_i)| \tag{11}$$

Putting (6) and (11) together, we get that

$$\textsc{cost}(\mathcal{G}^{i-1}, \mathcal{C}_i) + \Phi(\mathcal{G}^i) - \Phi(\mathcal{G}^{i-1}) \leq (\lambda + 9)\sqrt{|V_i|} + (8 - \frac{\lambda}{2}) \sum_{j=1}^{h} s_j + (3 - \lambda)|\mathcal{B}(\mathcal{G}^{i-1}, \mathcal{C}_i)|$$

By letting λ be a sufficiently large constant (namely $\lambda = 16$), we get that

$$\textsc{cost}(\mathcal{G}^{i-1}, \mathcal{C}_i) + \Phi(\mathcal{G}^i) - \Phi(\mathcal{G}^{i-1}) = O(\sqrt{|V_i|}). \qquad \square$$

Corollary 3.6. *Let G be a 3-regular plane graph with ν vertices. For a sequence $\mathscr{C} = \mathcal{C}_1, \ldots, \mathcal{C}_N$ of flarbable fleeqs for graph $G = (V, E)$ where $\nu = |V|$,*

$$\sum_{i=1}^{N} \textsc{cost}(\mathcal{G}^{i-1}, \mathcal{C}_i) = O(\nu + N\sqrt{\nu + N})$$

Proof. Using the result of Theorem 3.5, we can write

$$\sum_{i=1}^{N} \textsc{cost}(\mathcal{G}^{i-1}, \mathcal{C}_i) + \Phi(\mathcal{G}^N) - \Phi(G) = O(N\sqrt{|V_i|}).$$

Because $\Phi(G) = \lambda \sum_{f \in F} \mu(f)$, we know that $\Phi(G) = O(\nu)$. Analogously, since each flarb operation adds at most 2 vertices by Lemma 2.2, we know that the number of vertices in $\mathcal{G}^N$ is $O(\nu + N)$ which, in turn, implies that $\Phi(\mathcal{G}^N) = O(\nu + N)$. Therefore,

$$\sum_{i=1}^{N} \textsc{cost}(\mathcal{G}^{i-1}, \mathcal{C}_i) = O(N\sqrt{|V_i|} + \Phi(G) - \Phi(\mathcal{G}^N)) = O(\nu + N\sqrt{\nu + N}) \qquad \square$$

4 The lower bound

In Section 4, we present an example of a 3-regular Halin graph G with ν vertices—a tree with all leaves connected by a cycle to make it 3-regular—and a corresponding flarb operation with cost $\Omega(\sqrt{\nu})$ that yields a graph isomorphic to G. Because this sequence can be repeated, the amortized cost of a flarb is $\Theta(\sqrt{\nu})$.

Let $\nu = 2k(k+1) - 2$ for some positive integer k. The construction of the 3-regular graph with ν vertices is depicted in Figure 3. In this graph, we show the existence of a flarbable curve $\mathcal{C}$ (dashed in the figure) such that the flarb operation on G produces a graph $\mathcal{G}(G, \mathcal{C})$ isomorphic to G. Moreover, $\mathcal{C}$ crosses at least k augmented $\mathcal{C}$-faces and k shrinking $\mathcal{C}$-faces. Therefore, $\textsc{cost}(G, \mathcal{C}) \geq k = \Omega(\sqrt{\nu})$ by Lemma 2.6. Since the resulting graph is isomorphic to the original graph, this operation can be repeated in succession an arbitrarily high number times. That is, there is a sequence of N flarbable curves $\mathcal{C}_1, \ldots, \mathcal{C}_N$ such that $\sum_{i=1}^{N} \textsc{cost}(\mathcal{G}^{i-1}, \mathcal{C}_i) = \Omega(N\sqrt{\nu})$,

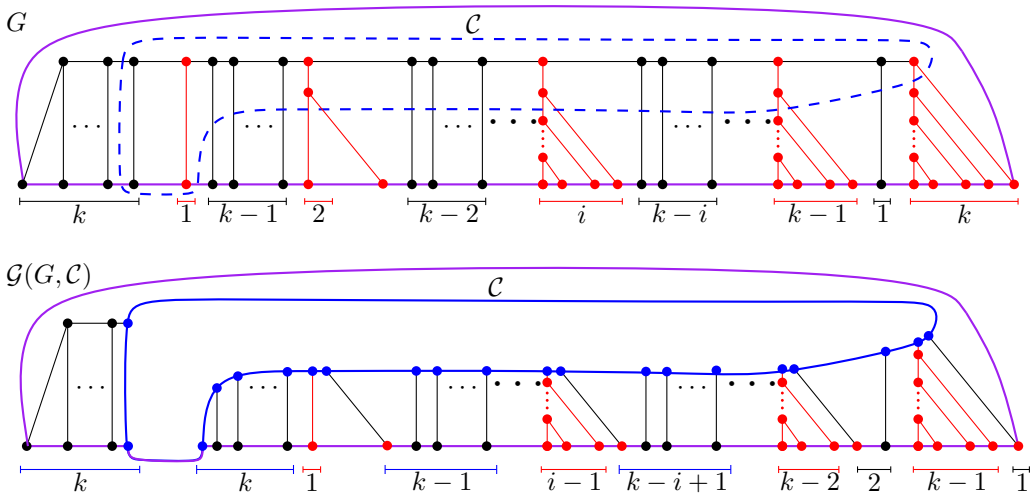

Figure 3: A 3-regular graph G with $\nu = 2k(k+1) - 2$ vertices. A flarbable curve $\mathcal{C}$ induces a flarb such that $\mathcal{G}(G, \mathcal{C})$ is isomorphic with G.

5 Computing the flarb

In this section, we describe a data structure to maintain the Voronoi diagram of a set S of n sites in convex position as new sites are added to S. Our structure allows us to find the edges of each preserved face and ignore them, thereby focusing only on necessary modifications to the combinatorial structure. The time we spend in these operations is then proportional to the number of non-preserved edges. Since this number is

proportional to the cost of the flarb, our data structure supports site insertions in time that is almost optimal (up to a polylogarithmic factor).

5.1 Grappa trees

Grappa trees [2] are a modification of link-cut trees, a data structure introduced by Sleator and Tarjan [15] to maintain the combinatorial structure of trees. They support the creation of new isolated vertices, the *link* operation which adds an edge between two vertices in disjoint trees, and the *cut* operation which removes an edge, splitting a tree into two trees.

We use this structure to maintain the combinatorial structure of the incrementally constructed Voronoi diagram $\mathcal{V}(S)$ of a set S of sites in convex position throughout construction. Recall that each insertion defines a flarbable curve $\mathcal{C}$, namely the boundary of the Voronoi cell of the inserted site. Our algorithm performs this flarb operation in time $O(\textsc{cost}(\mathcal{V}(S), \mathcal{C}) \log^7 n)$, where n is the number of vertices inserted so far. That is, we obtain an algorithm whose running time depends on the minimum number of link and cut operations that the Voronoi diagram, which is a tree, must undergo after each insertion. Moreover, this Voronoi diagram answers nearest neighbor queries in $O(\log n)$ time.

A grappa tree, as introduced by Aronov et al. [2], is a data structure is based on the worst-case version of the link-cut tree construction of Sleator and Tarjan [15]. This structure maintains a forest of fixed-topology trees subject to many operations, including Make-Tree, Link, and Cut, each in $O(\log n)$ worst-case time while using $O(n)$ space.

As in [2, 15], we decompose a rooted binary tree into a set of maximal vertex-disjoint downward paths, called *heavy paths*, connected by tree edges called *light edges*. Each heavy path is in turn represented by a biased binary tree whose leaf-nodes correspond to the vertices of the heavy path. Non-leaf nodes represent edges of this heavy path, ordered in the biased tree according to their depth along the path. Therefore, vertices that are higher (closer to the root) in the path correspond to leaves farther left in the biased tree. Each leaf node ℓ of a biased tree B represents an internal vertex v of the tree which has a unique light edge l_v adjacent to it. We keep a pointer from ℓ to this light edge. Note that the other endpoint of l_v is the root of another heavy path which in turn is represented by another biased tree, say B'. We merge these two biased trees by adding a pointer from ℓ to the root of B'. After merging all the biased trees in this way, we obtain the grappa tree of a tree T. A node of the grappa tree that is an internal vertex of its biased tree represents a heavy edge and has two children, whereas a node that is a leaf of its biased tree represents a vertex of the heavy path (and its unique adjacent light edge) and has only one child. By a suitable choice of paths and biasing, as described in [15], the grappa tree has height $O(\log n)$.

In addition, grappa trees allow us to store left and right marks on each of its nodes, i.e., on each edge of T. To assign the mark of a node, grappa trees support the $O(\log n)$-time operation Left-Mark(T, v, m_l) which sets the mark m_l to every edge in the path from v to the super root of T (Right-Mark(T, v, m_l) is defined analogously). In our setting, we use the marks of an edge e to keep track of the faces adjacent to this edge in a geometric embedding of T. Since T is rooted, we can differentiate between the left and the right faces adjacent to e.

The following definition formalizes the operations supported by a grappa tree.

Definition 5.1. Grappa trees solve the following data-structural problem: maintain a forest of rooted binary trees with specified topology subject to:

T = Make-Tree(v): Create a new tree T with a single internal vertex v (not previously in another tree).

T = Link(v, w): Given a vertex v in one tree T_v and the root w of a different tree T_w, connect v and w and merge T_v with T_w into a new tree T.

(T_1, T_2) = Cut(e): Delete the existing edge $e = (v, w)$ in tree T, splitting into T two trees T_1 and T_2 containing v and w, respectively.

Evert(v): Make external node v the root of its tree, reversing the orientation (which endpoint is closer to the root) of every edge along the root-to-v path.

Left-Mark(T, v, m_ℓ): Set the left mark of every edge on the root-to-v path in T to the new mark m_ℓ, overwriting the previous left marks of these edges.

Right-Mark(T, v, m_r): Set the right mark of every edge on the root-to-v path in T to the new mark m_r, overwriting the previous right marks of these edges.

$(e, m_\ell^*, m_r^*) =$ Oracle-Search(T, O_e): Search for the edge e in tree T. The data structure can find e only via *oracle queries*: given two incident edges f and f' in T, the provided oracle $O_e(f, f', m_\ell, m_r, m'_\ell, m'_r)$ determines in constant time which "side" of f contains e, i.e., whether e is in the component of $T - f$ that contains f', or in the rest of the tree (which includes f itself). The data structure provides the oracle with the left mark m_ℓ and the right mark m_r of edge f, as well as the left mark m'_ℓ and the right mark m'_r of edge f', and at the end, it returns the left mark m_ℓ^* and the right mark m_r^* of the found edge e.

Theorem 5.2. *[Theorem 7 from [2]] A grappa tree maintains the combinatorial structure of a forest and supports each operation described above in $O(\log n)$ worst-case time per operation, where n is the total size of the trees affected by the operation.*

5.2 The Voronoi diagram

Let S be a set of n sites in convex position and let $\mathcal{V}(S)$ be the binary tree representing the Voronoi diagram of S. We store $\mathcal{V}(S)$ using a grappa tree. In addition, we assume that each edge of $\mathcal{V}(S)$ has two *face-markers*: its left and right markers which store the sites of S whose Voronoi regions are adjacent to this edge on the left and right side, respectively. While a grappa tree stores only the topological structure of $\mathcal{V}(S)$, with the aid of the face-markers we can retrieve the geometric representation of $\mathcal{V}(S)$. Namely, for each vertex v of $\mathcal{V}(S)$, we can look at its adjacent edges and their face-markers to retrieve the point in the plane representing the location of v in the Voronoi diagram of S in $O(1)$ time. Therefore, we refer to v also as a point in the plane. Recall that each vertex v of $\mathcal{V}(S)$ is the center of a circle that passes through at least three sites of S, we call these sites the *definers* of v and we call this circle the *definer circle* of v.

Observation 5.3. *Given a new site q in the plane such that $S' = S \cup \{q\}$ is in convex position, the vertices of $\mathcal{V}(S)$ that are closer to q than to any other point of S' are exactly the vertices whose definer circle encloses q.*

Let q be a new site such that $S' = S \cup \{q\}$ is in convex position. Let $\textsc{cell}(q, S')$ be the Voronoi region of q in the Voronoi diagram of S' and let $\partial\textsc{cell}(q, S')$ denote its boundary. Recall that we can think of $\mathcal{V}(S)$ as a Halin graph by connecting all its leaves by a cycle to make it 3-regular. While we do not explicitly use this cycle, we need it to make our definitions consistent. In this Halin graph, the curve $\partial\textsc{cell}(q, S')$ can be made into a closed curve by going around the leaf of $\mathcal{V}(S)$ contained in $\textsc{cell}(q, S')$, namely the point at infinity of the bisector between the two neighbors of q along the convex hull of S'. In this way, $\partial\textsc{cell}(q, S')$ becomes a flarbable curve. Therefore, we are interested in performing the flarb operation it induces which leads into a transformation of $\mathcal{V}(S)$ into $\mathcal{V}(S')$.

5.3 Heavy paths in Voronoi diagrams

Recall that for the grappa tree of $\mathcal{V}(S)$, we computed a heavy path decomposition of $\mathcal{V}(S)$. In this section, we first identify the portion of each of these heavy paths that lies inside $\textsc{cell}(q, S')$. Once this is done, we test if any edge adjacent to an endpoint of these paths is preserved. Then within each heavy path, we use the biased trees built on it to further find whether there are non-preserved edges on this heavy path. After identifying all the non-preserved edges, we remove them, which results in a split of $\mathcal{V}(S)$ into a forest where each edge in $\textsc{cell}(q, S')$ is preserved. Finally, we show how to link the disjoint components back to the tree resulting from the flarb operation.

We first find the heavy paths of $\mathcal{V}(S)$ whose roots lie in $\textsc{cell}(q, S')$. Additionally, we find the portion of each of these heavy paths that lies inside $\textsc{cell}(q, S')$.

Recall that there is a leaf ρ of $\mathcal{V}(S)$ that lies in $\textsc{cell}(q, S')$: the point at infinity of the bisector between

the two neighbors of q along the convex hull of S'. As a first step, we root $\mathcal{V}(S)$ at ρ by calling Evert(ρ). In this way, ρ becomes the root of $\mathcal{V}(S)$ and all the heavy paths have a *root* which is their endpoint closest to ρ.

Let R be the set the of roots of all heavy paths of $\mathcal{V}(S)$, and let $R_q = \{r \in R : r \in \text{CELL}(q, S')\}$. We focus now on computing the set R_q. By Observation 5.3, each root in R_q has a definer circle that contains q. We use a dynamic data structure that stores the definer circles of the roots in R and returns those circles containing a given query point efficiently.

Lemma 5.4. *There is a fully dynamic $O(n)$-space data structure to store a set of circles (not necessarily with equal radii) that can answer queries of the form: Given a point q in the plane, return a circle containing q, where insertions take $O(\log^3 n)$ amortized time, deletions take $O(\log^6 n)$ amortized time, and queries take $O(\log^2 n)$ worst-case time.*

Proof. Chan [6] presented a fully dynamic randomized data structure that can answer queries about the convex hull of a set of n points in three dimensions where insertions take $O(\log^3 n)$ amortized time, deletions take $O(\log^6 n)$ amortized time, and extreme-point queries take $O(\log^2 n)$ worst-case time. We use this structure to solve our problem, but first, we must transform our input into an instance that can be handled by this data structure.

Let $\mathscr{C}$ be the dynamic set of circles we want to store. Consider the paraboloid-lifting which maps every point $(x, y) \to (x, y, x^2 + y^2)$. Using this lifting, we identify each circle $C \in \mathscr{C}$ with a plane π_C in $\mathbb{R}^3$ whose intersection with the paraboloid projects down as C in the xy-plane. Moreover, a point $q = (x, y)$ lies inside C if and only if point $(x, y, x^2 + y^2)$ lies below the plane π_C.

Let $\Pi = \{\pi_C : C \in \mathscr{C}\}$ be the set of planes corresponding to the circles in $\mathscr{C}$. In the above setting, our query can be translated as follows: Given a point $q' = (x, y, x^2 + y^2)$ on the paraboloid, find a plane $\pi_C \in \Pi$ that lies above q'.

Using standard point-plane duality in $\mathbb{R}^3$, we can map the set of planes Π to a point set Π^*, and a query point q' to a plane q^* such that a query translates to a *plane query*: Given a query plane q^*, find a point of Π^* that lies below it.

Using the data structure introduced by Chan [6] to store Π^*, we can answer plane queries as follows. Consider the direction orthogonal to q^* pointing in the direction below q^*. Then, find the extreme point of the convex hull of Π^* in this direction in $O(\log^2 n)$ time. If this extreme point lies below q^*, return the circle of $\mathscr{C}$ corresponding to it. Otherwise, we return that no point of Π^* lies below q^*, which implies that no circle of $\mathscr{C}$ contains q. Insertions take $O(\log^3 n)$ time while removals from the structure take $O(\log^6 n)$ time. □

For our algorithm, we store each root in R into the data structure given by Lemma 5.4. Using this structure, we obtain the following result.

Lemma 5.5. *We can compute each root in R_q in total $O(|R_q| \log^6 n)$ amortized time.*

Proof. After querying for a root whose definer circle contains q, we remove it from the data structure and query it again to find another root with the same property until no such root exists. Since queries and removals take $O(\log^2 n)$ and $O(\log^6 n)$ time, respectively, we can find all roots in R_q in $O(|R_q| \log^6 n)$ time. □

Given a root $r \in R$, let h_r be the heavy path whose root is r. Because the portion of $\mathcal{V}(S)$ that lies inside $\text{CELL}(q, S')$ is a connected subtree, we know that, for each $r \in R_q$, the portion of the path h_r contained in $\text{CELL}(q, S')$ is also connected. In order to compute this connected subpath, we want to find the last vertex of h_r that lies inside of $\text{CELL}(q, S')$, or equivalently, the unique edge of h_r having exactly one endpoint in the interior of $\text{CELL}(q, S')$. We call such an edge the *q-transition edge* of h_r (or simply transition edge).

Lemma 5.6. *For a root $r \in R_q$, we can compute the transition edge of h_r in $O(\log n)$ time.*

Proof. Let e_r be the transition edge of h_r. We make use of the oracle search proper of a grappa-tree to find the edge e_r. To this end, we must provide the data structure with an oracle such that: given two incident edges f and f' in $\mathcal{V}(S)$, the oracle determines in constant time which side of f contains the edge e_r, i.e., whether e_r is in the component of $\mathcal{V}(S) \setminus f$ that contains f', or in the rest of the tree (which includes f itself).

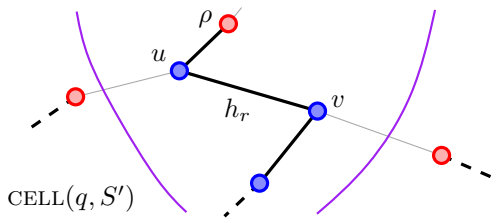

Figure 4: Path h_r contains two adjacent vertices u and v such that the light edge of u is a left edge while the light edge of v is a right edge. The edge uv cannot be preserved.

The data structure provides the oracle with the left and the right marks of f and f'. Given such an oracle, a grappa tree allows us to find the edge e_r in $O(\log n)$ time by Theorem 5.2.

Given two adjacent edges f and f' of $\mathcal{V}(S)$ that share a vertex v, we implement the oracle described above as follows. Recall that the left and right face-marks of f and f' correspond to the sites of S whose Voronoi region is incident to the edges f and f'. Thus, we can determine the definers of the vertex v, find their circumcircle, and test whether q lies inside it or not in constant time. Thus, by Observation 5.3, we can test in $O(1)$ time whether v lies in R_q or not and hence, decide if e_r is in the component of $\mathcal{V}(S) \setminus f$ that contains f', or in the rest of the tree. □

5.4 Finding non-preserved edges.

Observation 5.7. *Given a 3-regular graph G and a flarbable curve $\mathcal{C}$, if we can test whether a point is enclosed by $\mathcal{C}$ in $O(1)$ time, then we can test whether an edge is preserved in $O(1)$ time.*

Proof. First note that we can test in $O(1)$ time whether an edge reappears by testing whether its two adjacent edges are fleeq-edges. Since a preserved edge is either an edge that reappears or a fleeq-edge adjacent to an edge that reappears, this takes only $O(1)$ time. □

Let $\mathcal{V}_q(S)$ be the subtree induced by all the edges of $\mathcal{V}(S)$ that intersect CELL(q, S'). Now, we work towards showing how to identify each non-preserved edge of $\mathcal{V}_q(S)$ in the fleeq induced by ∂CELL(q, S'). For each root $r \in R_q$, we compute the transition edge e_r of h_r using Lemma 5.6 in $O(\log n)$ time per edge. Assume that w is the vertex of e_r that is closer to r (or is equal to r). We consider each edge adjacent to w and test whether it is preserved. Since each vertex of $\mathcal{V}_q(S)$ has access to its definers via the face markers of its incident edges, we can test if this vertex lies in CELL(q, S'). Thus, by Observation 5.7, we can decide whether an edge of $\mathcal{V}_q(S)$ is preserved in $O(1)$ time.

We mark each non-preserved edge among them as *shadow*. Because we can test whether an edge is preserved in $O(1)$ time, and since computing e_r takes $O(\log n)$ time by Lemma 5.6, this can be done in total amortized $O(|R_q| \log n)$ time. In addition, notice that if h_r contains two adjacent vertices u and v such that the light edge of u is a left edge while the light edge of v is a right edge (or vice versa), then the edge uv cannot be preserved; see Figure 4. In this case, we say that uv is a *bent edge*. We want to mark all the bent edges in $\mathcal{V}_q(S)$ as shadow, but first we need to identify them efficiently.

Note that it suffices to find all the bent edges of h_r for a given root $r \in R_q$, and then repeat this process for each root in R_q. To find the bent edges in h_r, we further extend the grappa tree in such a way that the biased tree representing h_r allows us to search for bent edges in $O(\log n)$ time. This extension is described as follows. Recall that each leaf s_v of a biased tree corresponds to a vertex v of the heavy path and has a pointer to the unique light edge adjacent to v. Since each light edge is either left or right, we can extend the biased tree to allow us to search in $O(\log n)$ time for the first two consecutive leaves where a change in direction occurs. From there, standard techniques allow us to find the next change in direction in additional $O(\log n)$ time. Therefore, we can find all the bent edges of a heavy path h_r in $O(\log n)$ time per bent edge. After finding each bent edge in h_r, we mark it is as a shadow edge.

Lemma 5.8. *An edge of $\mathcal{V}_q(S)$ is a preserved edge if and only if it was not marked as a shadow edge.*

Proof. Since we only mark non-preserved edges as shadow, we know that if an edge is preserved, then it is not shadow.

Assume that there is a non-preserved edge uv of $\mathcal{V}_q(S)$ that is not marked as shadow. If uv is a heavy edge, then it belongs to some heavy path h_r for some $r \in R_q$. We know that uv cannot be the transition edge of h_r since it would have been shadowed when we tested whether it was preserved. Thus, uv is completely contained in $\textsc{cell}(q, S')$. We can also assume that uv is not a bent edge, otherwise uv would have been shadowed. Therefore, the light children of u and v are either both left or both right children, say left. Since uv is not preserved, either the light child of u or the light child of v must be inside $\textsc{cell}(q, S')$. Otherwise if both edges cross the boundary of $\textsc{cell}(q, S')$, then uv is preserved by definition.

Assume that u has a light left child r' that is inside $\textsc{cell}(q, S')$. That is, r' must be the root of some heavy path and hence belongs to R_q. However, in this case we would have checked all the edges adjacent to u while processing the root $r' \in R_q$. Therefore, every edge that is non-shadow and intersects $\textsc{cell}(q, S')$ is a preserved edge. □

Corollary 5.9. *It holds that* $\sigma = \Theta(\textsc{cost}(\mathcal{V}(S), \partial\textsc{cell}(q, S')))$.

Let σ be the number of shadow edges of $\mathcal{V}(S)$, which is equal to the number of non-preserved edges by Lemma 5.8. The following relates the size of R_q with the value of σ.

Lemma 5.10. *It holds that* $|R_q| = O(\sigma \log n)$.

Proof. Given a root r of R_q, let p_r be the parent of r and notice that the edge rp_r is a light edge that is completely contained in R_q. Note that p_r belongs to another heavy path h_t, for some $t \in R_q$. If p_r is the endpoint of the transition edge of h_t closest to the root, then we add a *dependency pointer* from r to t. This produces a *dependency graph* with vertex set R_q. Since there is only transition edge per heavy path, that the in-degree of each vertex in this dependency graph is one. Therefore, the dependency graph is a collection of (oriented) *dependency paths*.

Since any path from a vertex to the root ρ of $\mathcal{V}(S)$ traverses $O(\log n)$ light edges, each dependency path has length $O(\log n)$. Let $r \in R_q$ be the sink of a dependency path. Consider the light edge rp_r and notice that it cannot be preserved, as p_r is not incident to a transition edge. Therefore, we can charge this non-preserved edge to the dependency path with sink r. Since a non-preserved can be charged only once, we have that σ is at least the number of dependency paths. Finally, as each dependency path has length $O(\log n)$, there are at least $\Omega(|R_q|/\log n)$ of them. Therefore $\sigma = \Omega(|R_q| \log n)$, or equivalently, $|R_q| = O(\sigma \log n)$ which yields our result. □

5.5 The compressed tree

Let $\mathcal{F}$ be the forest obtained from $\mathcal{V}_q(S)$ by removing all the shadow edges (this is just for analysis purposes, so far no cut has been performed). Note that each connected component of $\mathcal{F}$ consists only of preserved edges that intersect $\textsc{cell}(q, S')$. Thus, each component inside $\textsc{cell}(q, S')$ is a comb, with a path as a *spine* and each child of a spine vertex pointing to the same side; see Figure 5. Thus, we have right and left combs, depending on whether the children of the spine are left or right children.

Our objective in the long term is to cut all the shadow edges and link the remaining components in the appropriate order to complete the flarb. To this end, we would like to perform an Eulerian tour on the subtree $\mathcal{V}_q(S)$ to find the order in which the subtrees of $\mathcal{V}(S) \setminus \mathcal{V}_q(S)$ that hang from the leaves of $\mathcal{V}_q(S)$ appear along this tour. However, this may be too expensive as we want to perform this in time proportional to the number of shadow edges and the size of $\mathcal{V}_q(S)$ may be much larger. To make this process efficient, we compress $\mathcal{V}_q(S)$ by contracting each comb of $\mathcal{F}$ into a single super node. By performing an Eulerian tour around this compressed tree, we obtain the order in which each component needs to be attached. We construct the compressed flarb and then we decompress as follows.

Note that each comb has exactly two shadow edges that connect it with the rest of the tree. Thus, we contract the entire component containing the comb into a single *super node* and add a left or right dummy child to it depending on whether this comb was left or right, respectively; see Figure 5. After the compression,

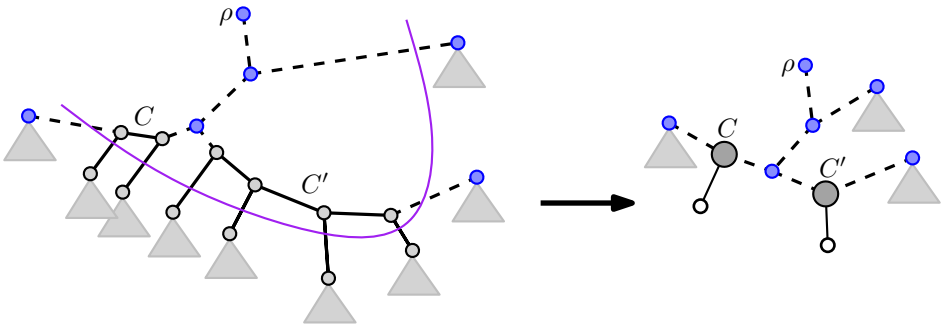

Figure 5: Two combs of $\mathcal{F}$ that are compressed into super nodes with their respective dummy leaves. An Eulerian tour around the compressed tree provides us with the order in which the trees hanging outside of $\textsc{cell}(q, S')$ should be attached.

the shadow edges together with the super nodes and the dummy vertices form a tree called the *compressed tree* that has $O(\sigma)$ vertices and edges, where σ is the total number of shadow edges.

Lemma 5.11. *We can obtain the compressed tree in $O(\sigma \log \sigma)$ time.*

Proof. Notice that each shadow edge is adjacent to two faces—its left face and its right face. Recall that each face ibounds the Voronoi cell of some site in S and that each shadow edge has two markers pointing to the sites defining its adjacent faces. Using hashing, we can group the shadow edges that are adjacent to the same face in $O(\sigma)$ time. Since preserved faces have no shadow edges on their boundary, we have at most $O(\sigma)$ groups.

Finally, we can sort the shadow edges adjacent to a given face along its boundary. To this end, we use the convex hull position of the sites defining the faces on the other side of each of these shadow edges. Computing this convex hull takes $O(\sigma \log \sigma)$ time. Once the shadow edges are sorted along a face, we can walk and check whether consecutive shadow edges are adjacent. If they are not, then the path between them consists only of preserved edges forming a comb; see Figure 5. Therefore, we can compress this comb and continue walking along the shadow edges. Since each preserved edge that reappears is adjacent to a face containing at least one shadow edge (namely the face that is not preserved), all the combs will be compressed during this procedure. □

The compressed tree is then a binary tree where each super node has degree three and each edge is a shadow edge. We now perform an Eulerian tour around this compressed tree and retrieve the order in which the leaves of this tree are visited. Some leaves are dummy leaves and some of them are original leaves of $\mathcal{V}_q(S)$; see Figure 5.

5.6 Completing the flarb

We now proceed to remove each of the shadow edges which results in a (compressed) forest with $O(\sigma)$ components. Note that each of the original leaves of $\mathcal{V}_q(S)$ was connected with its parent via a shadow edge and hence it lies now as a single component in the resulting forest. For each of these original leaves of $\mathcal{V}_q(S)$, we create a new *anchor* node and link it as the parent of this leaf. Moreover, there could be internal vertices that become isolated. In particular this will be the case of the root ρ. These vertices are deleted and ignored for the rest of the process. To complete the flarb, we create two new nodes ρ' and ρ'' which will be the two new leaves of the Voronoi diagram, one of them replacing ρ. Then, we construct a path with endpoints ρ and ρ' that connects the super nodes and the anchor nodes according to the traversal order of their leaves; see Figure 6. The resulting tree is a *super comb* Y, where each vertex on the spine is either a super node or an anchor node, and all the leaves are either dummy leaves or original leaves of $\mathcal{V}_q(S)$. Since we combined $O(\sigma)$ components into a tree, we need $O(\sigma)$ time.

We proceed now to decompress Y. To decompress a super node of Y that corresponds to a comb, we

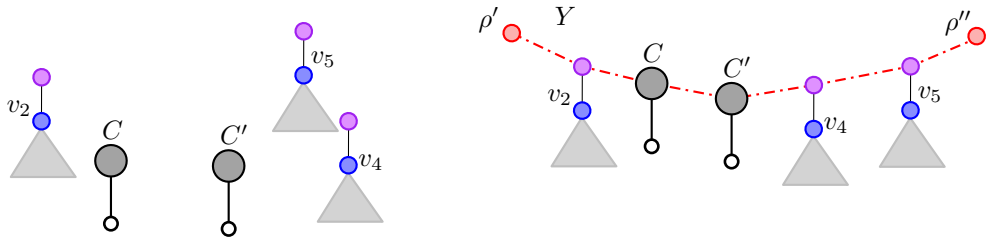

Figure 6: Left: An anchor node is created for each isolated leaf of $\mathcal{V}_q(S)$ and attached as its parent. Other isolated nodes are ignored. Right: A super comb is created connecting two new leaves ρ' and ρ'' through a path. This path connects anchor and super nodes in the order retrieved by the Eulerian tour around the compressed tree.

consider the two neighbors of the super node in Y and attach each of them to the ends of the spine of the comb. For an anchor node, we simply note that there is a component of $\mathcal{V}(S)$ hanging from its leaf; see Figure 7. In this way, we obtain all the edges that need to be linked. After the decompression, we end with the tree $\mathcal{V}(S')$ resulting from the flarb. Thus, the flarb operation of inserting q can be implemented with $O(\sigma)$ link and cuts.

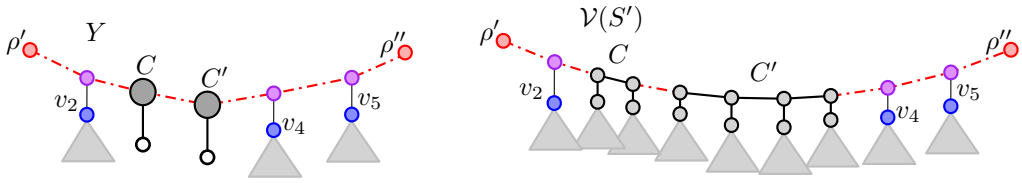

Figure 7: The tree $\mathcal{V}(S')$ achieved after the decompression.

Recall that any optimal algorithm needs to perform a cut for each edge that is not preserved. Since each non-preserved edge is shadow by Lemma 5.8, the optimal algorithm needs to perform at least $\Omega(\sigma)$ operations. Therefore, our algorithm is optimal and computes the flarb using $\Theta(\sigma)$ link and cuts. Moreover, by Lemmas 5.5 and 5.6 we can compute the flarb in $O(|R_q| \log^6 n + \sigma \log n)$ amortized time using $\Theta(\sigma)$ link and cuts. Since $|R_q| = O(\sigma \log n)$ by Lemma 5.10, we obtain the following.

Theorem 5.12. *The flarb operation of inserting q can be implemented with $O(K)$ link and cuts, where K is the cost of the flarb. Moreover, it can be implemented in $O(K \log^7 n)$ amortized time.*

References

[1] G. Aloupis, L. Barba, and S. Langerman. Circle separability queries in logarithmic time. In *Proceedings of the 24th Canadian Conference on Computational Geometry*, CCCG'12, pages 121–125, August 2012.

[2] B. Aronov, P. Bose, E. D. Demaine, J. Gudmundsson, J. Iacono, S. Langerman, and M. Smid. Data structures for halfplane proximity queries and incremental Voronoi diagrams. In *LATIN 2006: Theoretical Informatics*, pages 80–92. Springer, 2006.

[3] L. Barba. Disk constrained 1-center queries. In *Proceedings of the 24th Canadian Conference on Computational Geometry*, CCCG'12, pages 15–19, August 2012.

[4] J. L. Bentley and J. B. Saxe. Decomposable searching problems i. static-to-dynamic transformation. *Journal of Algorithms*, 1(4):301–358, 1980.

[5] P. Bose, S. Langerman, and S. Roy. Smallest enclosing circle centered on a query line segment. In *Proceedings of the 20th Canadian Conference on Computational Geometry (CCCG 2008)*, pages 167–170, 2008.

[6] T. M. Chan. A dynamic data structure for 3-d convex hulls and 2-d nearest neighbor queries. *Journal of the ACM (JACM)*, 57(3):16, 2010.

[7] Y.-J. Chiang and R. Tamassia. Dynamic algorithms in computational geometry. *Proceedings of the IEEE*, 80(9):1412–1434, 1992.

[8] M. De Berg, M. Van Kreveld, M. Overmars, and O. C. Schwarzkopf. *Computational geometry*. Springer, 2000.

[9] H. Edelsbrunner and R. Seidel. Voronoi diagrams and arrangements. *Discrete & Computational Geometry*, 1(1):25–44, 1986. ISSN 0179-5376.

[10] R. Klein. *Concrete and Abstract Voronoi Diagrams*, volume 400 of *Lecture Notes in Computer Science*. Springer, 1989. ISBN 3-540-52055-4.

[11] R. Klein, K. Mehlhorn, and S. Meiser. Randomized incremental construction of abstract voronoi diagrams. *Computational Geometry*, 3(3):157 – 184, 1993. ISSN 0925-7721.

[12] R. Klein, E. Langetepe, and Z. Nilforoushan. Abstract voronoi diagrams revisited. *Computational Geometry*, 42(9):885 – 902, 2009. ISSN 0925-7721.

[13] M. H. Overmars. *The design of dynamic data structures*, volume 156. Springer Science & Business Media, 1983.

[14] S. Pettie. Applications of forbidden 01 matrices to search tree and path compression-based data structures. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1457–1467, 2010.

[15] D. D. Sleator and R. E. Tarjan. A data structure for dynamic trees. *Journal of Computer and System Sciences*, 26(3):362–391, 1983.